arXiv:cs/0110022v1 [cs.PL] 9 Oct 2001

Mixed-Initiative Interaction = Mixed Computation*

Naren Ramakrishnan, Robert Capra, and Manuel A. Pérez-Quiñones
Department of Computer Science
Virginia Tech, Blacksburg, VA 24061, USA
Contact Email: naren@cs.vt.edu

February 1, 2008

Abstract

We show that partial evaluation can be usefully viewed as a programming model for realizing mixed-initiative functionality in interactive applications. Mixed-initiative interaction between two participants is one where the parties can take turns at any time to change and steer the flow of interaction. We concentrate on the facet of mixed-initiative referred to as 'unsolicited reporting' and demonstrate how out-of-turn interactions by users can be modeled by 'jumping ahead' to nested dialogs (via partial evaluation). Our approach permits the view of dialog management systems in terms of their native support for staging and simplifying interactions; we characterize three different voice-based interaction technologies using this viewpoint. In particular, we show that the built-in form interpretation algorithm (FIA) in the VoiceXML dialog management architecture is actually a (well disguised) combination of an interpreter and a partial evaluator.

*This work is supported in part by US National Science Foundation grants DGE-9553458 and IIS-9876167.

1 Introduction

Mixed-initiative interaction [HM97] has been studied for the past 30 years in the areas of artificial intelligence planning [VCP+95], human-computer interaction [BWFH92], and discourse analysis [Cou77]. As Novick and Sutton point out [NS97], it is 'one of those things that people think that they can recognize when they see it even if they can't define it.' It can be broadly viewed as a flexible interaction strategy between participants where the parties can take turns at any time to change and steer the flow of interaction. Human conversations are typically mixed-initiative and, interestingly, so are interactions with some modern computer systems. Consider the following two dialogs with a telephone pizza delivery service that has voice-recognition capability (the line numbers are provided for ease of reference):

Dialog 1

0 **Caller:** ≺calls Joe's Pizza on the phone≻
1 **System:** Thank you for calling Joe's pizza ordering system.
2 **System:** What size pizza would you like?
3 **Caller:** I'd like a medium, please.
4 **System:** What topping would you like on your pizza?
5 **Caller:** Pepperoni.
6 **System:** What type of crust do you want?
7 **Caller:** Uh, deep-dish.
8 **System:** So that is a medium pepperoni pizza with deep-dish crust. Is this correct?
9 **Caller:** Yes.
(conversation continues to get delivery and payment information)

Dialog 2

0 **Caller:** ≺calls Joe's Pizza on the phone≻
1 **System:** Thank you for calling Joe's pizza ordering system.
2 **System:** What size pizza would you like?
3 **Caller:** I'd like a sausage pizza, please.
4 **System:** Okay, sausage.
5 **System:** What size pizza would you like?
6 **Caller:** Medium.
7 **System:** What type of crust do you want?
8 **Caller:** Deep-dish.
9 **System:** So that is a medium sausage pizza with deep-dish crust. Is this correct?
10 **Caller:** Yes.
(conversation continues to get delivery and payment information)

Both these conversations involve the specification of a (size,topping,crust) tuple to complete the pizza ordering procedure but differ in important ways. In the first dialog, the caller responds to the questions in the order they are posed by the system. The system has the initiative at all times (other than, perhaps, on line 0) and such an interaction is thus referred to as *system-initiated*. In the second dialog, when the system prompts the caller about pizza size, he responds with information about his choice of topping instead (sausage; see line 3 of *Dialog 2*). Nevertheless, the conversation is not stalled and the system continues with the other aspects of the information gathering activity. In particular, the system registers that the caller has specified a topping, skips its default question on this topic, and repeats its question about the size (see line 5 of *Dialog 2*). The caller thus gained the initiative for a brief period during the conversation, before returning it to the system. A conversation that 'mixes' these modes of interaction in such arbitrary ways is said to be *mixed-initiative*.

1.1 Tiers of Mixed-Initiative Interaction

It is well accepted that mixed-initiative provides a more natural and personalized mode of interaction. A matter of debate, however, are the qualities that an interaction must possess to merit its classification as mixed-initiative [NS97]. In fact, determining who has the initiative at a given point in an interaction can itself be a contentious issue! The role of intention in an interaction and the underlying task goals also affect the characterization of initiative. We will not attempt to settle this debate here but a few preliminary observations will be useful.

One of the basic levels of mixed-initiative is referred to as *unsolicited reporting* in [AGH99] and is illustrated in *Dialog 2* above. In this facet, a participant provides information out-of-turn (in our case the caller, about his choice of topping). Furthermore, the out-of-turn interaction is not agreed upon in advance by the two participants. Novick and Sutton [NS97] stress that the unanticipated nature of out-of-turn interactions is important and that mere turn-taking (perhaps in a hardwired order) does not constitute mixed-initiative. Finally, notice that in *Dialog 2* there is a resumption of the original questioning task once processing of the unsolicited response is completed. In other applications, an unsolicited response might shift the control to a new interaction sequence and/or abort the current interaction.

Another level of mixed-initiative involves *subdialog invocation*; for instance, the computer system might not have understood the user's response and ask for clarifications (which amounts to it having the initiative). A final, sophisticated, form of mixed-initiative is one where participants negotiate with each other to determine initiative (as opposed to merely 'taking the initiative') [AGH99]:

Dialog 3

(with apologies to O. Henry)
Husband: Della, Something interesting happened today that I want to tell you.
Wife: I too have something exciting to tell you, Jim.
Husband: Do you want to go first or shall I tell you my story?

In addition to models that characterize initiative, there are models for designing dialog-based interaction systems. Allen et al. [ABD+01] provide a taxonomy of such software models — finite-state machines, slot-and-filler structures, frame-based methods, and more sophisticated models involving planning, agent-based programming, and exploiting contextual information. While mixed-initiative interaction can be studied in any of these models, it is beyond the scope of this paper to address all or even a majority of them.

Instead, we concentrate on the view of (i) a dialog as a task-oriented information assessment activity requiring the filling of a set of slots, (ii) where one of the participants in the dialog is a computer system and the other is a human, and (iii) where mixed-initiative arises from unsolicited reporting (by the human), involving out-of-turn interactions. This characterization includes many voice-based interfaces to information (our pizza ordering dialog is an example) and web sites modeling interaction by hyperlinks [RP01]. In Section 2, we show that partial evaluation can be usefully viewed as a programming model for such applications. Section 3 presents three different voice-based interaction technologies and analyzes them in terms of their native support for mixing initiative. Finally, Section 4 discusses other facets of mixed-initiative and mentions other software models to which our approach can be extended.

2 Programming a Mixed-Initiative Application

Before we outline the design of a system such as Joe's Pizza, we introduce a notation [Lev83, Gof76] that captures basic elements of initiative and response in an interaction sequence. The notation expresses the local organization of a dialog [PQn96, PQnS96] as adjacency pairs; for instance, *Dialog 1* is represented as:

(Ic Rs) (Is Rc) (Is Rc) (Is Rc) (Is Rc)
0 1 2 3 4 5 6 7 8 9

The line numbers given below the interaction sequence refer to the utterance numbers in the dialog presented in Section 1. The letter I denotes who has the initiative — caller (c) or the system (s) — and the letter R denotes who provides the response. It is easy to see from this notation that *Dialog 1* consists of five turns and that the system has the initiative for the last four turns. The initial turn is modeled as the caller having the initiative because he or she chose to place the phone call in the first place. The system quickly takes the initiative after playing a greeting to the caller (which is modeled here as the response to the caller's call). The subsequent four interactions then address three questions and a confirmation, all involving the system retaining the initiative (Is) and the caller in the responding mode (Rc). Likewise, the mixed-initiative interaction in *Dialog 2* is represented as:

(Ic Rs) (Is (Ic Rs) Rc) (Is Rc) (Is Rc)
0 1 2,5 3 4 6 7 8 9 10

In this case, the system takes the initiative in utterance 2 but instead of responding to the question of size in utterance 3, the caller takes the initiative, causing an 'insertion' to occur in the interaction sequence (dialog) [Lev83]. The system responds with an acknowledgement ('Okay, sausage.') in utterance 4. This is represented as the nested pair (Ic Rs) above. The system then re-focuses the dialog on the question of pizza size in utterance 5 (thus retaking the initiative). In utterance 6 the caller responds with the desired size (medium) and the interaction proceeds as before, from this point.

The notation is useful to describe the space of possible interactions that are to be supported. For instance, utterances 0 and 1 have to proceed in order. Utterances dealing with selection of (size,topping,crust) can then be nested in any order and provide interesting opportunities for mixing initiative. For instance, if a user is a frequent customer of Joe's Pizza, he might take the initiative and specify all three pizza attributes on the first available prompt:

Dialog 4

0 **Caller:** ≺calls Joe's Pizza on the phone≻
1 **System:** Thank you for calling Joe's pizza ordering system.
2 **System:** What size pizza would you like?
3 **Caller:** I'd like a sausage pizza, medium, and deep-dish.
(conversation continues with confirmation of order)

Finally, the utterances dealing with confirmation of the user's request can proceed only after choices of all three pizza attributes have been made. There are 13 possible interaction sequences (discounting permutations of attributes specified in a given utterance) — 1 possibility of specifying everything in one utterance, 6 possibilities of specification in two utterances, and 6 possibilities of specification in three utterances. If we include permutations, there are 24 possibilities (our calculations do not consider situations where, for instance, the system doesn't recognize the user's input and reprompts for information).

Many programming models view mixed-initiative sequences as requiring some special attention to be accommodated. In particular, they rely on recognizing when a user has provided unsolicited input[1] and qualify a shift-in-initiative as a 'transfer of control.' This implies that the mechanisms that handle out-of-turn interactions are often different from those that realize purely system-directed interactions. Fig. 1 (left) describes a typical software design. A dialog manager is in charge of prompting the user for input, queueing messages onto an output medium, event

[1]We use the term 'unsolicited input' here to refer to expected but out-of-turn inputs as opposed to completely unexpected (or out-of-vocabulary) inputs.

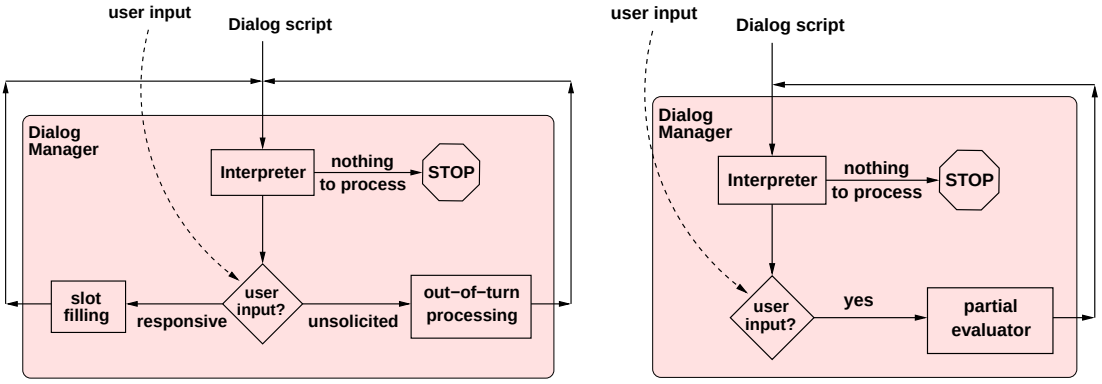

Figure 1: Designs of software systems for mixed-initiative interaction. (left) Traditional system architecture, distinguishing between responsive and unsolicited inputs. (right) Using partial evaluation to handle inputs uniformly.

processing, and managing the overall flow of interaction. One of its inputs is a dialog script that contains a specification of interaction and a set of slots that are to be filled. In our pizza example, slots correspond to placeholders for values of size, topping, and crust. An interpreter determines the first unfilled slot to be visited and presents any prompts for soliciting user input. A responsive input requires mere slot filling whereas unsolicited inputs would require out-of-turn processing (involving a combination of slot filling and simplification). In turn, this causes a revision of the dialog script. The interpreter terminates when there is nothing left to process in the script. While typical dialog managers perform miscellaneous functions such as error control, transferring to other scripts, and accessing scripts from a database, the architecture in Fig. 1 (left) focuses on the aspects most relevant to our presentation.

Our approach, on the other hand, is to think of a mixed-initiative dialog as a program, all of whose arguments are passed by reference and which correspond to the slots comprising information assessment. As usual, an interpreter in the dialog manager queues up any applicable prompts to an output medium. Both responsive and unsolicited inputs by a user now correspond (uniformly) to values for arguments; they are processed by partially evaluating the program with respect to these inputs. The result of partial evaluation is another dialog (simplified as a result of user input) which is handed back to the interpreter. This novel design is depicted in Fig. 1 (right) and a dialog script represented in a programmatic notation is given in Fig. 2. Partial evaluation of Fig. 2 with respect to user input will remove the conditionals for all slots that are filled by the utterance (global variables are assumed to be under the purview of the interpreter). The reader can verify that a sequence of such partial evaluations will indeed mimic the interaction sequence depicted in *Dialog 2* (and any of the other mixed-initiative sequences).

Partial evaluation serves two critical uses in our design. The first is obvious, namely the processing of out-of-turn interactions (and any appropriate simplifications to the dialog script). The more subtle advantage of partial evaluation is its support for staging mixed-initiative interactions. The mix-equation [Jon96, JGS93] holds for every possible way of splitting inputs into two categories, without enumerating and 'trapping' the ways in which the computations can be staged. For instance, the nested pair in *Dialog 2* is supported as a natural consequence of our design, not by anticipating and reacting to an out-of-turn input.

Another way to characterize the system designs in Fig. 1 is to say that Fig. 1 (left) makes a distinction of responsive versus unsolicited inputs, whereas Fig. 1 (right) makes a more fundamental distinction of fixed-initiative (interpretation) versus mixed-initiative (partial evaluation). In other words, Fig. 1 (right) carves up an interaction sequence into (i) turns that are to be handled in the order they are modeled (by an interpreter), and (ii) turns that can involve mixing of initiative (handled by a partial evaluator). In the latter case, the computations are actually used as a *representation of interactions.* Since only mixed-initiative interactions involve multiple staging options and since these are handled by the partial evaluator, our design requires the *least* amount of specification (to support all interaction sequences). For instance, the script in Fig. 2 models the parts that involve mixing of initiative and helps

```
pizzaorder(size,topping,crust)
{
  if (unfilled(size)){
   /* prompt for size */
  }
  if (unfilled(topping)){
   /* prompt for topping */
  }
  if (unfilled(crust)){
   /* prompt for crust */
  }
}
```

Figure 2: Modeling a dialog script as a program parameterized by slot variables that are passed by reference.

realize all of the 13 possible interaction sequences. At the same time it does not model the confirmation sequence of *Dialog 2* because the caller cannot confirm his order before selecting the three pizza attributes! This turn should be handled by straightforward interpretation.

To the best of our knowledge, such a model of partial evaluation for mixed-initiative interaction has not been proposed before. An extensive literature search has revealed no related prior work. While computational models for mixed-initiative interaction remain an active area of research [HM97], such work is characterized by keywords such as 'controlling mixed-initiative interaction,' 'knowledge representation and reasoning strategies,' and 'multi-agent co-ordination.' There are even projects that talk about 'integrating' mixed-initiative interaction and partial evaluation to realize an architecture for planning and learning [VCP+95]. We are optimistic that our work has the same historical significance as the relation between explanation-based generalization (a learning technique in AI) and partial evaluation established by van Haremelen and Bundy in 1988 [vHB88].

3 Software Technologies for Voice-Based Mixed-Initiative Applications

One of the main contributions of our model is that it characterizes the minimum amount of information needed to program a mixed-initiative interaction sequence. Once a programmer supplies a script such as Fig. 2 mixed-initiative interaction is obtained, quite literally, 'for free.' This means that we can use the design in Fig. 1 (right) as a benchmark to compare and contrast the amount of specification required in other approaches.

As indicated in Section 1.1, our model is applicable to voice-based interaction technologies as well as web access via hyperlinks. We concentrate on voice-based applications since interaction with web sites is addressed in a related paper [RP01] and because the design constraints in voice-based applications pose interesting considerations for our model. In addition, a variety of commercial technologies are available for voice-based applications (in contrast to web sites) that will aid in comparative assessment.

3.1 Basic Principles of Voice-Based Interaction

Before we can study the programming of mixed-initiative in a voice-based application, it will be helpful to understand the basic architecture (see Fig. 3) of a spoken language processing system. As a user speaks into the system, the sounds produced are captured by a microphone and converted into a digital signal by an analog-to-digital converter. In telephone-based systems (the VoiceXML architecture covered later in the paper is geared toward this mode), the microphone is part of the telephone handset and the analog-to-digital conversion is typically done by equipment in the telephone network (in some cellular telephony models, the conversion would be performed in the handset itself).

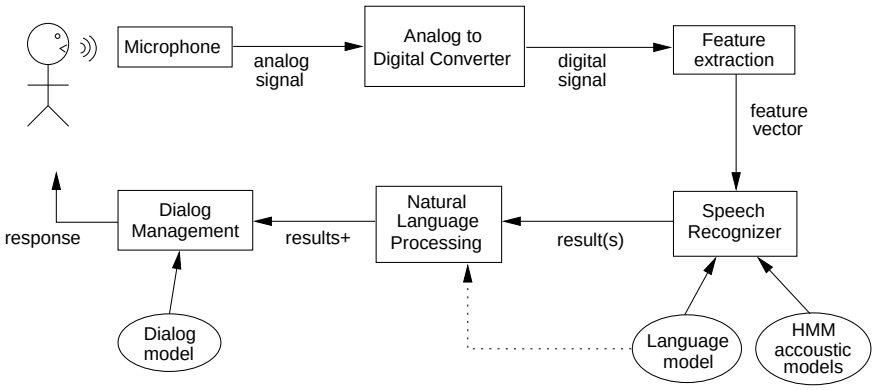

Figure 3: Basic components of a spoken language processing system.

The next stage (feature extraction) prepares the digital speech signal to be processed by the speech recognizer. Features of the signal important for speech recognition are extracted from the original signal, organized as an attribute vector, and passed to the recognizer.

Most modern speech recognizers use Hidden Markov Models (HMMs) and associated algorithms to represent, train, and recognize speech. HMMs are probabilistic models that must be trained on an input set of data. A common technique is to create sets of acoustic HMMs that model phonetic units of speech in context. These models are created from a training set of speech data that is (hopefully) representative of the population of users who will use the system. A language model is also created prior to performing recognition. The language model is typically used to specify valid combinations of the HMMs at a word- or sentence-level. In this way, the language model specifies the words, phrases, and sentences that the recognizer can attempt to recognize. The process of recognizing a new input speech signal is then accomplished using efficient search algorithms (such as Viterbi decoding) to find the best matching HMMs, given the constraints of the language model. The output of the speech recognizer can take several different forms, but the basic result is a text string that is the recognizer's best guess of what the user said. Many recognizers can provide additional information such as a lattice of results, or an N-best ranked list of results (in case the later stages of processing wish to reject the recognizer's top choice). A good introduction to speech recognition is available in [JM00].

The stages after speech recognition vary depending on the application and the types of processing required. Fig. 3 presents two additional phases that are commonly included in spoken language processing systems in one form or another. We will broadly refer to the first post-recognition stage as natural language processing (NLP). NLP is a large field in its own right and includes many sub-areas such as parsing, semantic interpretation, knowledge representation, and speech acts; an excellent introduction is available in Allen's classic [All95]. Our presentation in this paper has assumed NLP support for slot-filling (i.e., determining values for slot variables from user input).

This is commonly achieved by defining parts of a language model and associating them with slots. The language model could take two major forms — context-free grammars and statistical-based (such as n-grams). Here we focus on the former: in this approach, slots can be specified within the productions of a context-free grammar (akin to a attribute grammar) or they can be associated with the non-terminals in the grammar.

We will refer to the next phase of processing as simply 'dialog management' (see Fig. 3). In this phase, augmented results from the NLP stage are incorporated into the dialog and any associated logic of the application is executed. The job of the dialog manager is to track the proceedings of the dialog and to generate appropriate responses. This is often done within some logical processing framework and a dialog model (in our case, a dialog script) is supplied as input that is specific to the particular application being designed. The execution of the logic on the dialog model (script) results in a response that can be presented back to the user. Sometimes response generation

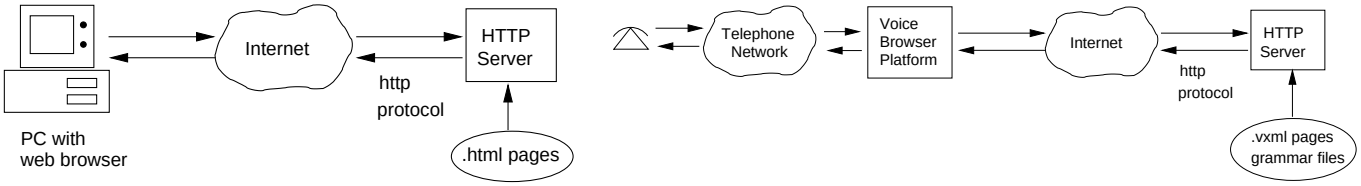

Figure 4: (left) Accessing HTML documents via a HTTP web server. (right) Accessing VoiceXML documents via a HTTP web server.

is separated out into a subsequent stage.

3.2 The VoiceXML Dialog Management Architecture

There are many technologies and delivery mechanisms available for implementing Fig. 3's basic components. A popular implementation can be seen in the VoiceXML dialog management architecture. VoiceXML is a markup language designed to simplify the construction of voice-response applications [voi00]. Initiated by a committee comprising AT&T, IBM, Lucent Technologies, and Motorola, it has emerged as a standard in telephone-based voice user interfaces and in delivering web content via voice. We will hence cover this architecture in detail.

The basic idea is to describe interaction sequences using a markup notation in a VoiceXML *document.* As the VoiceXML specification [voi00] indicates, a VoiceXML document constitutes a conversational finite state machine and describes a sequence of interactions (both fixed- and mixed-initiative are supported). A web server can serve VoiceXML documents using the HTTP protocol (Fig. 4 (right)), just as easily as HTML documents are currently served over the Internet (Fig. 4 (left)). In addition, voice-based applications require a suitable delivery platform, illustrated by a telephone in Fig. 4 (right). The voice-browser platform in Fig. 4 (right) includes the VoiceXML interpreter which processes the documents, monitors user inputs, streams messages, and performs other functions expected of a dialog management system. Besides the VoiceXML interpreter, the voice-browser platform includes speech recognizers, speech synthesizers, and telephony interfaces to help realize important aspects of voice-based interaction.

Dialog specification in a VoiceXML document involves organizing a sequence of *forms* and *menus*. Forms specify a set of slots (called field item variables) that are to be filled by user input. Menus are syntactic shorthands (much like a `case` construct); since they involve only one field item variable (argument), there are no opportunities for mixing initiative. We do not discuss menus further in this paper. An example VoiceXML document for our pizza application is given in Fig. 5.

As shown in Fig. 5, the pizza dialog consists of two forms. The first form (`welcome`) merely welcomes the user and transitions to the second. The `place_order` form involves four `field`s (slot variables) — the first three cover the pizza attributes and the fourth models the confirmation variable (recall the dialogs in Section 1). In particular, prompts for soliciting user input in each of the fields are specified in Fig. 5.

Interactions in a VoiceXML application proceed just like a web application except instead of clicking on a hyperlink (to goto a new state), the user talks into a microphone. The VoiceXML interpreter then determines the next state to transition to. Any appropriate responses (to user input) and prompts are delivered over a speaker. The core of the interpreter is a so-called form interpretation algorithm (FIA) that drives the interaction. In Fig. 5, the fields provide for a fixed-initiative, system-directed interaction. The FIA simply visits all fields in the order they are presented in the document. Once all fields are filled, a check is made to ensure that the confirmation was successful; if not, the fields are cleared (notice the `clear namelist` tag) and the FIA will proceed to `prompt` for the inputs again, starting from the first unfilled field — `size`.

The form in Fig. 5 is referred to as a directed one since the computer has the initiative at all times and the `field`s are filled in a strictly sequential order. To make the interaction mixed-initiative (with respect to `size`, `crust`, and

```
<?xml version="1.0"?>
<vxml version="1.0">
<!-- pizza.vxml
     A simple pizza ordering demo to illustrate some basic elements
     of VoiceXML.  Several details have been omitted from this demo
     to help make the basic ideas stand out. -->
  <form id="welcome">
    <block name="block1">
      <prompt> Thank you for calling Joe's pizza ordering system. </prompt>
      <goto next="#place_order" />
    </block>
  </form>

  <form id="place_order">
    <field name="size">
      <prompt> What size pizza would you like? </prompt>
    </field>

    <field name="topping">
      <prompt> What topping would you like on your pizza? </prompt>
    </field>

    <field name="crust">
      <prompt> What type of crust do you want? </prompt>
    </field>

    <field name="verify">
      <prompt>
        So that is a <value expr="size"/> <value expr="topping"/> pizza
        with <value expr="crust"/> crust.
        Is this correct?
      </prompt>
      <grammar> yes | no </grammar>
    </field>

    <filled>
      <if cond="verify=='no'">
         <clear namelist="size topping verify crust"/>
         <prompt> Sorry.  Your order has been canceled. </prompt>
      <else/>
         <prompt>Thank you for ordering from Joe's pizza.</prompt>
      </if>
    </filled>

  </form>
</vxml>
```

Figure 5: Modeling the pizza ordering dialog in a VoiceXML document.

```
#JSGF V1.0;

grammar sizetoppingcrust;

public <sizetoppingcrust> =
   <size> {this.size=$} [<topping> {this.topping=$}] [<crust> {this.crust=$}] |
   <size> {this.size=$} <crust> {this.crust=$} <topping> {this.topping=$} |
   <topping> {this.topping=$} [<crust> {this.crust=$}] [<size> {this.size=$}] |
   <topping> {this.topping=$} <size> {this.size=$} <crust> {this.crust=$} |
   <crust> {this.crust=$} [<size> {this.size=$}] [<topping> {this.topping=$}] |
   <crust> {this.crust=$} <topping> {this.topping=$} <size> {this.size=$};

<size> = small | medium | large;
<topping> =  sausage | pepperoni | onions | green peppers;
<crust> = regular | deep dish | thin;
```

Figure 6: A form-level grammar to be used in conjunction with the script in Fig. 5 to realize mixed-initiative interaction. The above productions for `sizetoppingcrust` cover all possibilities of filling slot variables from user input, including multiple slots filled by a given utterance, and various permutations of specifying pizza attributes.

```
While (true)
{
        // SELECT PHASE
        Select the first form item with an unsatisfied guard condition
             (e.g., unfilled)
          If no such form item, exit

        // COLLECT PHASE
        Queue up any prompts for the form item
        Get an utterance from the user

        // PROCESS PHASE
        foreach (slot in user's utterance)
        {
               if (slot corresponds to a field item) {
                    copy slot values into field item variables
                    set field item's 'just_filled' flag
               }
        }
        // some code for executing any 'filled' actions triggered
}
```

Figure 7: Outline of the form interpretation algorithm (FIA) in the VoiceXML dialog management architecture. Adapted from [voi00].

```
#JSGF V1.0;

grammar sizetoppingcrust;

public <sizetoppingcrust> = word*;

word = <size> {this.size=$} |
       <crust> {this.crust=$} |
       <topping> {this.topping=$};

<size> = small | medium | large;
<topping> =  sausage | pepperoni | onions | green peppers;
<crust> = regular | deep dish | thin;
```

Figure 8: A alternative form-level grammar to realize mixed-initiative interaction with the script in Fig. 5.

`topping`), the programmer merely has to specify a so-called *form-level grammar* that describes possibilities for slot-filling from a user utterance. An example form-level grammar file (`sizetoppingcrust.gram`) that covers all possibilities is given in Fig. 6. The grammar is associated with the dialog script by including the line:

```
<grammar src="sizetoppingcrust.gram" type="application/x-jsgf"/>
```

just before the definition of the first `field` (size) in Fig. 5.

The form-level grammar contains productions for the various choices available for size, topping, and crust and also qualifies all possible parses for a given utterance (modeled by the non-terminal `sizetoppingcrust`). Any valid combination of the three pizza aspects uttered by the user (in any order) is recognized and the appropriate slot variables are instantiated. To see why this also achieves mixed-initiative, let us consider the FIA in more detail.

Fig. 7 only reproduces the salient aspects of the FIA relevant for our discussion. Compare the basic elements of the FIA to the stages in Fig. 1 (right). The Select phase corresponds to the interpreter, the Collect phase gathers the user input, and actions taken in the Process phase mimic the partial evaluator. Recall that 'programs' (scripts) in VoiceXML can be modeled by finite-state machines, hence the mechanics of partial evaluation are considerably simplified and just amount to filling the slot and removing it from further consideration. Since the FIA repeatedly executes till there are no remaining form items, the processing phase (Process) is effectively parameterized by the form-level grammar file in Fig. 6. In other words, the form-level grammar file not only enables slot filling, *it also implicitly directs the staging of interactions for mixed-initiative.* When the user specifies 'peperroni medium' in an utterace, not only does the grammar file enable the recognition of the slots they correspond to (topping and size), it also directs the FIA to simplify these slots (and remove them in any subsequent interaction).

The form-level grammar file shown in Fig. 6 (which is also a specification of interaction staging) may make VoiceXML's design appear overly complex. In reality, however, we could have used the vanilla form-level grammar file in Fig. 8. While helping to realize mixed-initiative with Fig. 5, the new form-level file (as does our model) also allows the possibility of utterances such as 'pepperoni pepperoni,' or even, 'pepperoni sausage!' Suitable semantics for such situations (including the role of side-effects) can be defined and accommodated in both the VoiceXML model and ours. It should thus be obvious to the reader that VoiceXML's dialog management architecture is actually implementing a mixed evaluation model (for conversational finite state machines), comprising interpretation and partial evaluation.

The VoiceXML specification [voi00] refers to the form-level file as merely a 'grammar file,' when it is actually also a specification of staging. Even though the grammar file serves the role of a language model in a voice application, we believe that separating its two functionalities is important in understanding mixed-initiative system

Software Technology	Support for Slot Simplification	Support for Interaction Staging
VoiceXML	$\surd$	$\surd$
Slot Filling Systems	$\surd$	$\times$
Recognizer-Only APIs	$\times$	$\times$

Table 1: Comparison of software technologies for voice-based mixed-initiative applications.

design. A case in point is our study of personalizing interaction with web sites [RP01]. There is no requirement for a 'grammar file,' as there is usually no ambiguity about user clicks and typed-in keywords. In this context, the functionality provided by our model is actually unmatched by any existing web-based interaction system (as web interfaces are not typically designed for mixing initiative). A way to incorporate mixed-initiative interaction into an existing interaction at a web site is described in [RP01].

3.3 Other Implementation Technologies

VoiceXML's FIA thus includes native support for slot filling, slot simplification, and interaction staging. All of these are functions enabled by partial evaluation in our model. Table 1 contrasts two other implementation approaches in terms of these aspects. In a purely slot-filling system, native support is provided for simplifying slots from user utterances but extra code needs to be written to model the control logic (for instance, 'the user still didn't specify his choice of size, so the question for size should be repeated.'). Several commercial speech recognition vendors provide APIs that operate at this level. In addition, many vendors support low-level APIs that provide basic access to recognition results (i.e., text strings) but do not perform any additional processing. We refer to these as recognizer-only APIs. They serve more as raw speech recognition engines and require significant programming to first implement a slot-filling engine and, later, control logic to mimic all possible opportunities for staging. Examples of the two latter technologies can be seen in the commercial spoken dialog systems market (from companies such as Nuance, IBM, and AT&T). The study presented in this paper suggests a systematic way by which their capabilities for mixed-initiative interaction can be assessed.

4 Discussion

Our work makes contributions to both partial evaluation and mixed-initiative interaction. For the partial evaluation community, we have identified a novel application where the motivation is the staging of interaction (rather than speedup). Since programs (dialogs) are used as specifications of interaction, they are *written to be partially evaluated*; partial evaluation is hence not an 'afterthought' or an optimization. A program can thus be thought of as a compaction of all possible interaction sequences that involve mixing initiative. An interesting research issue is: Given (i) a set of interaction sequences, and (ii) addressable information (such as arguments and slot variables), determine (iii) the smallest program so that every interaction sequence can be staged in the model of Fig. 1 (right). This requires algorithms to automatically decompose parts of interaction sequences into those that are best addressed in the interpreter and those that can benefit from representation and specialization by the partial evaluator.

For mixed-initiative interaction, we have presented a programming model that accommodates all possibilities of staging, without explicit enumeration. The model makes a distinction between fixed-initiative (and which has to be explicitly programmed) and mixed-initiative (specifications of which can be elegantly compressed for subsequent partial evaluation). We have identified instantiations of this model in VoiceXML and slot-filling APIs. We hope this observation will help system designers gain additional insight into voice application design strategies.

It should be recalled that there are various facets of mixed-initiative that are not addressed in this paper. Extending our programming model to cover these facets is an immediate direction of future research. For example,

VoiceXML's design currently supports dialogs such as the following:

Dialog 5

1 **System:** Thank you for calling Joe's pizza ordering system.
2 **System:** What size pizza would you like?
3 **Caller 1:** What sizes do you have?
3 **Caller 2:** Err.. Why don't you ask me the questions in topping-crust-size order?

Caller 1's request, while demonstrating initiative, implies a dialog with an optional stage (which cannot be modeled by partial evaluation). Such a situation has to be trapped by the interpreter, not by partial evaluation. *Caller 2* does specify a staging, but his staging poses constraints on the computer's initiative, not his own. Such a 'meta-dialog' facet [BWFH92] requires the ability to jump out of the current dialog; VoiceXML provides many elements for describing such transitions.

VoiceXML also provides certain 'impure' features and side-effects in its programming model. For instance, after selecting a size (say, medium), the caller could retake the initiative in a different part of the dialog and select a size again (this time, large). This will cause the new value to over-ride any existing value in the `size` slot (see Fig. 7). In our model, this implies the dynamic substitution of an earlier, 'evaluated out,' stage with a functional equivalent. Obviously, the dialog manager has to maintain some state (across partial evaluations) to accomplish this feature. We plan to investigate programming models suitable for these aspects. In addition, we plan to extend our software model beyond slot-and-filler structures, to include reasoning and exploiting context.

Our long-term goal is to characterize mixed-initiative facets, not in terms of initiative, interaction, or task models but in terms of the opportunities for staging and the program transformation techniques that can support such staging. This means that we can establish a taxonomy of mixed-initiative facets based on the transformation techniques (e.g., partial evaluation, slicing) needed to realize them. Such a taxonomy would also help connect the facets to design models for interactive software systems.

References

[ABD+01] J.F. Allen, D.K. Byron, M. Dzikovska, G. Ferguson, L. Galescu, and A. Stent. Towards Conversational Human-Computer Interaction. *AI Magazine*, 2001. to appear.

[AGH99] J.F. Allen, C.I. Guinn, and E. Horvitz. Mixed-Initiative Interaction. *IEEE Intelligent Systems*, Vol. 14(5):pages 14–23, Sep-Oct 1999.

[All95] J.F. Allen. *Natural Language Understanding*. Benjamin Cummings, 1995. Second Edition.

[BWFH92] H. Brunner, G. Whittemore, K. Ferrara, and J. Hsu. An Assessment of Written/Interaction Dialogue for Information Retrieval Applications. *Human-Computer Interaction*, Vol. 7:pages 197–249, 1992.

[Cou77] M. Coulthard. *An Introduction to Discourse Analysis*. Longman, London, 1977.

[Gof76] E. Goffman. Replies and Responses. *Language in Society*, Vol. 5:pages 257–313, 1976.

[HM97] S. Haller and S. McRoy. *Computational Models for Mixed Initiative Interactions (Papers from the 1997 AAAI Spring Symposium)*. Technical Report: SS-97-04: AAAI/MIT Press, 1997.

[JGS93] N.D. Jones, C.K. Gomard, and P. Sestoft. *Partial Evaluation and Automatic Program Generation*. Prentice Hall International, 1993.

[JM00] D. Jurafsky and J.H. Martin. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*. Prentice Hall, 2000.

[Jon96] N.D. Jones. An Introduction to Partial Evaluation. *ACM Computing Surveys*, Vol. 28(3):pages 480–503, September 1996.

[Lev83] S.C. Levinson. *Pragmatics*. Cambridge University Press, 1983. Cambridge Textbooks in Linguistics.

[NS97] D.G. Novick and S. Sutton. What is Mixed-Initiative Interaction? In S. Haller and S. McRoy, editors, *Procedings of the AAAI Spring Symposium on Computational Models for Mixed Initiative Interaction*, pages 114–116. AAAI/MIT Press, 1997.

[PQn96] M.A. Pérez-Quiñones. *Conversational Collaboration in User-Initiated Interruption and Cancellation Requests*. PhD thesis, The George Washington University, Washington DC, 1996.

[PQnS96] M.A. Pérez-Quiñones and J.L. Sibert. A Collaborative Model of Feedback in Human-Computer Interaction. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI'96)*, pages 316–323. Vancouver, BC, Canada, 1996.

[RP01] N. Ramakrishnan and S. Perugini. The Partial Evaluation Approach to Information Personalization. *ACM Transactions on Information Systems*, August 2001. Communicated for publication. Also available as Technical Report cs.IR/0108003, Computing Research Repository (CoRR) at `http://xxx.lanl.gov/abs/cs.IR/0108003`.

[VCP+95] M. Veloso, J. Carbonell, A. Pérez, D. Borrajo, E. Fink, and J. Blythe. Integrating Planning and Learning: The PRODIGY Architecture. *Journal of Experimental and Theoretical Artificial Intelligence*, Vol. 7(1):pages 81–120, 1995.

[vHB88] F. van Harmelen and A. Bundy. Explanation-Based Generalisation = Partial Evaluation. *Artificial Intelligence*, Vol. 36(3):pages 401–412, 1988.

[voi00] Voice eXtensible Markup Language: VoiceXML. Technical report, VoiceXML Forum, March 2000. Version 1.00.